

**WIRELESS INFRARED NETWORK TRANSCEIVER
AND METHODS AND SYSTEMS FOR OPERATING SAME**

Field of the Invention

- 5 The present invention relates generally to wireless networks and more particularly to a wireless mesh network incorporating infrared network nodes and methods and systems for aligning same.

Background of the Invention

- 10 In many areas of the world the World Wide Web (WWW), or Internet, has become a significant medium for the exchange of information including everything from casual electronic mail (e-mail) to legal and business documents to entertainment media. Much of the material exchanged over the Internet comprises very large electronic files, for example large documents, music, video
15 and even full-length motion pictures are available for exchange and distribution over the Internet.

- While commercial services often choose fast but expensive high-speed Internet connections for business purposes, typical consumer connections comprise
20 relatively slow telephone modems. For example, a typical commercial T1 connection will yield in the range of 1,544 kilobits per second (Kbps) or 1.544 megabits per second (Mbps) data communications rate at a monthly cost in the range of \$1,000 to \$2,000. In contrast, a typical consumer telephone modem connection will provide a 56 Kbps data communications rate at a cost of in the
25 range of \$10-\$30 month.

- As commercial services provide richer content for consumer use, data file sizes increase. For example, a typical audio music file may be in the range of 3-5 Megabytes and take up to 10 minutes for a consumer to download over a
30 telephone modem. A typical audio/video file, for example a full-length movie, may run in the thousands of mega bytes size range and take a significant part of a day for a consumer to download over a modem. Streaming audio or video may require the download of large quantities of data over an extended period of time.

It is obvious that the ability of commercial services to provide rich, large media files is rapidly outstripping the typical consumer's ability to receive those files.

5 Recently, several affordable, high-speed alternatives have become available to the traditional consumer telephone modem. Cable modems use the cable television infrastructure to provide Internet connections having a speed of about 1,100 Kbps, about 20x times the speed of a telephone modem. DSL modems use conventional telephone lines to provide Interconnect connections, and have an average speed of about 700 Kbps, or over 12x times the speed of a telephone
10 modem. Both cable and DSL modems are priced at approximately twice the cost of telephone modem services, with slightly higher equipment costs than for standard modems.

The higher speed cable and DSL connections are geographically limited,
15 however, by the underlying infrastructure. Many areas of the United States and worldwide include regions not serviced by cable television or where the cable television networks have not and will not be upgraded to support high-speed data modems. Similarly, DSL service is not available in many geographic areas. Numerous reasons exist for the limited availability of cable and DSL services,
20 including high cost of infrastructure upgrade, technological limitations, physical geographical limitations and, in some areas, low demand. As with many types of commercial services, the incremental costs of extending infrastructure are becoming increasingly higher, sometimes by multiples or even exponentially, as attempts are made to expand those infrastructures to every last consumer.

25

There thus exists a real demand for high-speed Internet connections in areas that cable and/or DSL service providers may never serve. This demand will increase as more content is provided and more business is executed over the Internet.

30 Some providers have attempted to expand service coverage while avoiding the high costs associated with expanding network infrastructure. This is typically accomplished using a wireless network, for example extending from an access point in the wired infrastructure. Wireless networks may be installed without the need for the wired infrastructure.

One type of wireless network uses wireless radio frequency (RF) components that transmit data in the radio frequency spectrum. These networks, however, have the disadvantages of being expensive and relatively slow. In one
5 embodiment of multi-channel multipoint distribution system (MMDS), for example, broadcasts occur at speeds up to 25Mbps but require very expensive spread-spectrum infrastructure equipment, costing on average \$16 million. MMDS customer equipment is also very expensive, with the cost of deployment for a single customer running in the range of \$1,000.

10 Wi-Fi, or 802.11b is a much lower speed technology; achieving throughputs of up to 5.5 Mbps full duplex. Wi-Fi bandwidth decreases significantly with distance between components, and is particularly dependent on obstructions such as roofs or walls, as well as interference from other networks or even microwave
15 ovens.

Another type of wireless network uses light, in the form of, for example, lasers or light-emitting diodes (LEDs) to transmit high-speed data in a process called free space optic systems, or FSO systems. While FSO systems are a cost-
20 effective high-speed communications medium, they require very highly aligned line-of-sight paths. More specifically, existing free space optic systems have very narrow beam divergence parameters requiring precision alignment. For this reason, laser and FSO components tend to be expensive and require high levels of maintenance and service.

25 There thus exists demand for high-speed, affordable Internet connections in geographies and neighborhoods into which more traditional, wired high-speed network infrastructure cannot be cost-effectively extended. This demand will grow significantly as the Internet is increasingly used to deliver content,
30 facilitate business transactions and support other matters amenable to electronic data transfer.

Summary of the Invention

The present invention uses networked, wireless infrared nodes to cost-effectively provide high-speed data capacity, including Internet access, to selected geographies and neighborhoods. While not thus limited, the network of the present invention can be cost-effectively extended to many areas not supporting traditional wired network infrastructure.

In accordance with one embodiment of the invention there is provided an infrared transceiver node, comprising: a support assembly; at least two infrared transceivers connected to the support assembly, each of the infrared transceivers including an infrared transmitter including a light-emitting diode and a transmitter lens for generating an infrared beam having a beam divergence of in the range of about 0.5 to 1 degree, an infrared receiver including a receiver lens and a photo detector for receiving an infrared beam and converting the infrared beam to an electronic signal, a first drive assembly for adjusting the elevation of the infrared transceiver, a second drive assembly for adjusting the azimuth of the infrared transceiver, and a processor connected to each of the first and second drive assemblies for controlling the azimuth and elevation of the infrared transceiver, a connector for connecting the infrared transceiver node to a user computer; and a switching engine connected to the connector and to each of the at least two infrared transceivers for switching data between the user computer and the at least two infrared transceivers.

In accordance with another embodiment of the invention there is provided a network of infrared transceiver nodes, comprising: a network controller; a plurality of infrared transceiver nodes, each of the infrared transceiver nodes including a support assembly; at least two infrared transceivers connected to the support assembly, each of the infrared transceivers including an infrared transmitter including a light-emitting diode and a transmitter lens for generating an infrared beam having a beam divergence of not less than about 0.5-1.0 degrees, an infrared receiver including a receiver lens and a photo detector for receiving an infrared beam and converting the infrared beam to an electronic signal, a first drive assembly for adjusting the elevation of the infrared transceiver, a second drive assembly for adjusting the azimuth of the infrared

transceiver, and a processor connected to each of the first and second drive assemblies for controlling the azimuth and elevation of the infrared transceiver; a connector for connecting the infrared transceiver node to a user computer; a switching engine connected to the connector and to each of the at least two

5 infrared transceivers for switching data between the user computer and the at least two infrared transceivers; each of the plurality of infrared transceiver nodes positioned outdoors on the surface of a support structure; the plurality of infrared transceiver nodes relatively positioned so that each infrared transceiver node has a line-of-sight to at least one other infrared transceiver; and means for

10 connecting at least one of the plurality of infrared transceiver nodes to the network controller.

In accordance with another embodiment of the invention there is provided a point-to-point transceiver, comprising: a transmitter for generating a beam

15 having a beam divergence of in the range of about 0.5 to 1 degree; a receiver having a visible window smaller than the beam divergence for receiving a beam; a first drive assembly for adjusting the elevation of the transceiver; a second drive assembly for adjusting the azimuth of the transceiver; a connector for connecting the transceiver to a user computer; and a processor connected to each

20 of the first and second drive assemblies for controlling the azimuth and elevation of the infrared transceiver.

In accordance with another embodiment of the invention there is provided a transimpedance amplifier circuit, comprising: a first differential amplifier; a

25 second differential amplifier having its inputs connected to the outputs of the first differential amplifier; a third differential amplifier having its inputs connected to the outputs of the second differential amplifier; a first resistor connected between the input of the first differential amplifier and the output of the second differential amplifier; and a second resistor connected between the

30 input of the first differential amplifier and the output of the third differential amplifier.

In accordance with yet another embodiment of the invention there is provided a system and method of installing a new point-to-point communications

transceiver in a network of existing point-to-point communications transceivers, the method comprising the steps of: positioning a new transceiver in a line of sight to at least one existing transceiver; positioning the new transceiver to point in a predetermined compass direction; at a predetermined time, the new

5 transceiver performing a first sweep over a first predetermined range to detect a signal from the existing transceiver; at the predetermined time synchronously with the new transceiver, the existing transceiver performing a second sweep over a second predetermined range to detect a signal from the new transceiver; upon the completion of the first sweep of the new transceiver, the new

10 transceiver returning to a position of maximum detected signal strength; and upon the completion of the second sweep of the existing transceiver, the existing transceiver returning to a position of maximum detected signal strength.

In accordance with another embodiment of the invention there is provided a

15 system and method of maintaining the alignment of first and second communicating transceivers in a network of existing point-to-point communications transceivers, the method comprising the steps of: monitoring the signal strength received by the first transceiver; monitoring the signal strength received by the second transceiver; if either of the signal strength

20 received by the first transceiver or the signal strength received by the second transceiver falls below a predetermined threshold, then with the second transceiver stationary, tracking the first transceiver over a first predetermined path to determine a position of maximum signal strength detected by the first transceiver in the first predetermined path, after the completion of the tracking of

25 the first transceiver, returning the first transceiver to the position of maximum signal strength detected by the first transceiver in the first predetermined path, with the first transceiver stationary, tracking the second transceiver over a second predetermined path to determine a position of maximum signal strength detected by the second transceiver in the second predetermined path, and after

30 the completion of the tracking of the second transceiver, returning the second transceiver to the position of maximum signal strength detected by the second transceiver in the second predetermined tracking path.

Description of the Drawing Figures

These and other objects, features and advantages of the invention will become apparent from a consideration of the Detailed Description of the Invention when read in conjunction with the drawing figures, in which:

- 5 Figure 1 is a diagrammatic view of an infrared wireless mesh network in accordance with the invention;
- Figure 2A is a perspective view of a network node in accordance with the present invention;
- Figures 2B and 2C are front and rear perspective views, respectively, of an
10 infrared transceiver from a network node;
- Figure 3A is a perspective view of the network node without the cover;
- Figure 3B is a perspective view of the central shaft of the network node;
- Figure 3C is a perspective view of the mechanical coupler for coupling the infrared transceiver to the central shaft;
- 15 Figures 4A and 4B are side and front plan views of the infrared transceiver;
- Figure 4C is a cut view of the infrared transceiver along axis A-A of Figure 4A;
- Figure 4D is an enlarged partial view of infrared transmitter elements of Figure 4C;
- Figure 4E is an enlarged partial view of infrared receiver elements of Figure 4C;
- 20 Figure 5 is a diagrammatic view of the optical elements of the infrared optics showing the relative positioning of the receiver and transmitter optics in the transceiver casing;
- Figure 6 is a diagrammatic view showing the optical infrared beam spread of the transmit optics of Figure 5 at a first distance;
- 25 Figure 7 is a diagrammatic view showing the optical infrared beam spread of the transmit optics of Figure 5 at a second distance;
- Figure 8A is a diagrammatic view showing the optical infrared beam characteristics of the infrared receiver;
- Figure 8B is an enlarged diagrammatic view of the dome lens of Figure 8A
30 showing the optical infrared beam characteristics;
- Figure 9 is a diagrammatic view of the network node showing the interconnection of the infrared transceivers with the main system board and access connector;

- Figure 10 is a block diagram view of the network node system board from Figure 9 including interconnections to off-board components;
- Figure 11 is a block diagram view of the access device of Figures 9 and 10;
- Figure 12 is a block diagram view of a network node transceiver board including
- 5 interconnections to off-board components;
- Figure 13 is a block diagram showing the connection of selected components of Figure 12;
- Figure 14A is a block diagram of the transimpedance amplifier of Figure 12;
- Figure 14B is a schematic circuit diagram of the transimpedance amplifier of
- 10 Figure 14A;
- Figures 15A and 15B together show a process for installing a new node into an existing network;
- Figures 16A and 16B together show a process for the auto-acquisition of a newly installed node into an existing network;
- 15 Figure 17 shows a process for asynchronously initiating and executing a transceiver tracking process;
- Figure 18 shows a process for initiating and executing an unscheduled transceiver tracking process; and
- Figure 19 shows a process for asynchronously aligning a transceiver based on
- 20 digital signal quality.

Detailed Description of the Invention

- With reference now to the drawing Figures, embodiments of the invention are
- 25 shown and described wherein like elements are indicated by like reference numerals throughout.

Network Node

- With reference first to Figure 1, a diagram of a residential neighborhood 100 is
- 30 shown including multiple houses indicated at 102A-E. As is typical in neighborhoods, trees or other natural obstacles exist, indicated at 104A&B, which may block lines of sight between adjoining houses. A road 105 transects neighborhood 100 in a conventional manner. In accordance with the present invention, a series of infrared network nodes, indicated at 106A-E, are

positioned on the upper, outer surfaces of the homes, for example the roofs. It will be seen that house 102D has elected not to receive services from the described network and so no network node is positioned on that house.

- 5 Network node 106E is positioned on a base station 108, typically comprising a neighboring building or cell tower with access to a traditional wired network. In one embodiment of the invention, the base station 108 comprises a network system server 202 or a set of network of servers. Network system server(s) 202 provides network control and management features in a manner described below.
- 10 Network system server may also provide any other network services to each node, such as entertainment and caching services. The traditional wired network connects the base station 108 to the Internet through network system server(s) 202. In another embodiment of the invention, the wired network connection provides an access to a remote facility or facilities comprising one or more
- 15 network system servers and a connection(s) to the Internet.

As is described in further detail below, other communications paths such as telephone dial-ups or local RF networks can be used for communication of various control and alignment data during local network set-up and maintenance

20 processes.

Within neighborhood 100, various lines of sight between adjacent network nodes are indicated by dotted lines 110A-E.

- 25 In operation, network nodes 106A-D are configured in a mesh network configuration. High-speed digital data, in the form of Internet Protocol (IP) packet data, is transmitted between adjoining nodes, in the manner described below. As in a conventional IP packet-switching network, each network node examines the IP data and makes a routing decision based on the IP parameters.
- 30 The present embodiment is illustrated as connected to the Internet, but the invention is not thus limited and is equally applicable, for example, to implement a private local or wide area network. There is thus provided a high-speed, digital data, IP packet-switching network using cost-effective, flexibly positioned network nodes 106A-D. While the network nodes have been illustrated in a

mesh network relationship, it will be understood that they would similarly function in alternate network structures, including point-to-point networks and others. It will further be understood that while network nodes are specified to use IP protocol for data packet switching, they would similarly function with
5 alternate networking protocol like ATM.

Referring now to Figure 2A, B and C, there is shown an exemplary network node 106 containing four mechanically and electrically connected infrared transceivers, indicated at 108A-D. Each transceiver is seen to include a
10 respective transmitter lens 109A-D, receiver lens 111A-D, elevation motor 107A-D and azimuth motor 113A-D. Each transceiver includes a mount 115A-D for supporting the transceiver optics, and a base 117A-D about which the mount with optics can rotate both horizontally and vertically. As described in further detail below, in each transceiver the motors are connected between the
15 optics mount and the transceiver base for rotating the mount with optics relative to the base. A weatherproof housing 112 encloses the infrared transceivers, providing weatherproof access for electrical connections thereto in a manner described in further detail below. Weatherproof housing 112 includes a top 112A, a bottom 112B and a plastic tube or sleeve 112C (shown cut away)
20 enclosing the internal devices, the plastic sleeve selected from a material transparent to infrared light. An exemplary sole infrared transceiver 108A is illustrated in Figures 2B and 2C.

It will be understood that, excepting for position, nodes 106A-E are substantially
25 identical, as are the transceivers 108A-D within the nodes. Thus, any description of the mechanical, electrical and/or optical structure of a particular node or transceiver is applicable to substantially all nodes and transceivers.

Referring now to Figure 3A, network node 106 is shown with weatherproof
30 housing 112 removed. Infrared transceivers 108A-D can be seen to be generally concentrically mounted on a central shaft 114, shown in Figure 3B, with each transceiver fixed to the shaft by a locking clasp 116 best shown in Figure 3C. Upon initial fabrication, locking clasps 116 are used to position infrared transceivers 108A-D in known relative positions to facilitate the installation and

tracking process described below. In one embodiment of the invention central shaft 114 is hollow for supporting electronic power and data cables.

It will be understood that, due to the nature of each of the figures described
5 above, not every feature of each transceiver is visible in every figure.

With reference now to Figures 4A, B, C, D and E, additional features of exemplary transceiver 108A are seen. More particularly Figure 4D shows a light-emitting diode (LED) 122A positioned to emit light through a diffuser
10 sheet 120A for transmission by transmitter lens 109A. LED 122A preferably contains a reflector positioned behind its active area for directing more emitted light forward through diffuser sheet 120A.

Figure 4E shows a detector dome lens 116A positioned to partially surround an
15 avalanche photodiode (APD) 118A for receiving light from receiver lens 111A. As will be further described below, the transmit optics including LED 122A, diffuser sheet 120A and transmitter lens 109A, and the receive optics including receiver lens 111A, dome lens 116A and APD 118A are carefully selected such to optimize the performance of both the individual infrared transceivers and the
20 entirety of the mesh network.

Elevation motors 107A-D constitute conventional stepper motors controlled by electrical signals described below and positioned between each transceiver base and optical system mount so as to enable the individual rotation of each
25 transceiver's transmit/receive optics in a vertical plane. Similarly, azimuth motors 113A-D also constitute conventional stepper motors controlled by electrical signals described below and positioned between each transceiver base and optical system so as to enable the individual rotation of each transceivers transmit/receive optics in a horizontal plane. Each of transceivers 108A-D thus
30 has the capacity to position its transmitter and receiver optics completely independently of the remaining transceivers within node 106. It will be understood that conventional gear mechanisms may be used in conjunction with the elevation and azimuth motors so as to enable the motors to step in consecutively desired angular increments.

In one embodiment of the invention, the elevation and azimuth stepper motors are provided as 1.8-degree stepper motors operating at 3.3 volts with nominal 1 amp of current draw. The drive system is selected to provide +/- 15 degree
5 travel in elevation and +/- 178 degree travel in azimuth. A plastic, 0.25" pitch diameter pinion attached to the motor shaft drives a 3.5" pitch diameter plastic spur gear. A 14:1 (3.5/.25) gear ratio thus provides pointing resolution of 0.13degrees (1.8/14).- In another embodiment of the invention, pointing resolution can be further increased by half-stepping the motors or by using 0.9°
10 stepping motors. It will be understood that numerous other motor/gear embodiments will function to perform the desired results.

As will be described in further detail below, electrical signals for controlling the operation of each transceiver, including the respective elevation and azimuth
15 motors, are contained on the electronics boards associated with each network node system board and transceiver boards. Such control signals may be loaded onto the board at the factory and/or downloaded remotely from a system server once the node is physically positioned for use.

20 Transceiver Optics

With reference now to Figures 5-8, one embodiment of the transmitter and receiver optics for transceivers 108A-D is shown and described. The optics are identical in each of the transceivers. For purposes of illustration, they are described with respect to transceiver 108A.

25 Referring first to Figure 5, the relative positioning of the receive optics 140 and transmit optics 142 is shown, with the axes of LED 122A and APD 118A seen to be spaced 3.25 inches apart. The rear surface of transmitter lens 109A is spaced 4.31 inches from the front surface of diffuser 120A, while the rear surface of
30 receiver lens 111A is spaced 3.57 inches from the front surface of dome lens 116A. The transceiver mount 115A supports the transmit and receive components and is used to establish these dimensions.

With reference now to Figures 6 and 7, the transmit optics 142 is selected to provide a 0.8 degree beam divergence, shown at 150, such that at 30 meters from LED 122A the beam diameter is 0.4 meters (Figure 6), while at 400 meters from the LED the beam diameter is 5.6 meters (Figure 7). In one embodiment, these characteristics can be obtained by using the relative positioning dimensions shown in Figures 5 and 6, and by selecting the following components having the following characteristics:

- LED 122A comprises an Osram part # SFH 4301 having a wavelength of 950nm, a standard 3mm dome lens LED package, an angular divergence of +/- 10 degrees and an active area of 300 microns.
- Diffuser sheet 120A comprises a Physical Optics Corporation part #LSD5PE4-2, comprising polyester having a thickness of 0.004 inches and a Full-Width Half-Maximum angle (FWHM) of 5 degrees.
- Transmitter lens 109A is a plano-convex lens manufactured from Acrylic/Polycarbonate having a focal length of 125mm and a diameter of 50mm.

With reference now to Figures 8A-B, one embodiment of the receiving optics can be seen where receiving lens 111A and dome lens 116A are selected and relatively positioned (Figure 8A) so as to provide a 0.5mm beam dispersement 152 (Figure 8B).

In one embodiment, the receiver optical components include:

- APD by Hamamatsu, part # S2382, having a TO48 package, an FOV of 170 degrees and an active area of 500 microns
- Dome lens 116A of aspheric shape and acrylic material having a focal length of 5mm and a diameter of 7mm
- Receive lens 111A of aspheric shape and acrylic material and having a focal length of 63mm and a diameter of 100mm

A long-pass filtering dye or coating may be added to one or more of the optical elements in the receive path, for example dome lens 116A or receive lens 111A, to reduce background light of a wavelength lower than the transmission wavelengths. This will reduce the amount of visible light falling on the

detector. The insensitivity of the detector to mid- and far-infrared wavelengths obviates the need to filter longer wavelengths of background light.

5 With these exemplary parameters and components, the received beam will be focused at APD 118A. 100% of the field coverage over 100% of the desired aperture is maintained, with the spot size at the detector being much smaller than the active area of the APD such that substantially all of the received light is detected by the APD. With a temperature change of 25 degrees centigrade, the incoming beam will still be focused to a spot that is smaller than the active area
10 of APD 118A, thereby accommodating operating temperature-caused variations.

Thus sufficient energy is collected to operate APD 118A, both with the beam ideally positioned and with the beam shifted in accordance with the expected maximum variances.

15

In selecting the various parameters for the transmit and receive optics described above, it is anticipated that the distance between adjoining nodes will be on average 0.25 miles or less. It is thus desired to provide a transmit beam divergence sufficiently wide to enable transceivers in adjoining nodes to easily
20 reach and maintain alignment as described below, but yet with sufficient power to reliably transmit and receive infrared signals.

A transmit beam divergence in the range of 0.5-1.0 degrees, nominally 0.8 degrees, is sufficient to provide the desired operating characteristics. This
25 relatively wide beam divergence is counter to the extremely narrow beam divergence generally used for optical communications.

The visible field of view of APD 118A, is selected to be slightly smaller than the transmit beam divergence. In this manner, if one transceiver can receive light
30 and/or data from another transceiver, then the assumption can be made that the transmitting transceiver can also receive signals from the receiving transmitter. This asynchronous alignment ability, described in further detail below, is useful in establishing and maintaining alignment between communicating nodes.

It will be appreciated that the dispersement and related characteristics of the received beam have been obtained through the use of a relatively small and inexpensive APD in combination with a relatively small and inexpensive dome lens.

5

The present inventors have achieved the desired operating characteristics using affordable, components, particularly plastic lenses, plastic housings, off-the-shelf LEDs, and off-the-shelf APDs, thus making each network node affordable and readily constructed.

10

It will now be apparent that numerous other configurations of transmit and receive optics may be used to achieve the same functional results.

Network System & Electronics

15 With reference now to Figures 9 and 10, Figure 9 shows a diagrammatic view of network node 106 including an internally contained system board 170 connected to an external access device 172. Each transceiver 108A-D is shown diagrammatically mounted on the central shaft 114 and has associated with it a respective transceiver board 220A-D. A power supply board 188 further resides
20 within the weatherproof housing of network node 106. As will be shown in further detail below, system board 170 is connected to each respective transceiver board 220D by means of a cable. In one embodiment, various electrical and power conductors may extend through the hollow center of central shaft 114.

25

With reference to Figure 10, a block diagram of system board 170 in network node 106 is shown including a central processing unit 173 connected to a memory storage device 176, the memory storage device including appropriate combinations of magnetic, optical and semiconductor storage. As shown,
30 memory device 176 includes DRAM memory 176A and flash memory 176B for storing various program instructions and data as described below.

System board 170 further includes a Fast Ethernet/IP switching engine 182 connected to its own dedicated memory device 184. A bank 186 of four physical

layer (PHY) devices 186A-D is situated on system board 170, each PHY serving as a packet data interface between a transceiver and switching engine 182. In Figure 10, the PHY devices are illustrated as a bank of four Fast Ethernet (100-Base FX) PHYs 186A-D.

5

It will be understood that the signals between the switching engine and the PHY correspond to Fast Ethernet format generated by a media access controller (MAC), which, in the described embodiment, is integrated into the CPU and the switching engine. It will be understood that the MAC can comprise a separate component. It will be understood that other physical layer protocol and devices may be used as interfaces between the transceivers and the switching engine. In an alternate embodiment the PHY devices may also be integrated as part of the switching engine and the CPU.

15 Continuing with Figure 10, switching engine 182 is connected to CPU 173 by means of two interfaces. One of the interfaces is through a bridge device 185. Another interface is through a pair of Fast Ethernet PHYs, 180 and 181, connected back to back. Bridge interface 185 is used to exchange switch control and management information between switching engine 182 and CPU 173. The
20 bridge interface is also used for the exchange of management related IP data packets between the network and the CPU through the switching engine. As described in detail below, switching engine 182 functions to control the routing of high-speed network data.

25 It will be understood that in alternate embodiments the pair of PHYs, 180 and 181 can be eliminated and substituted, for example, by a compatible media-independent-interface provided by the CPU and switching engine. In another embodiment, either or each of the PHYs could be integrated with the CPU and/or the switching engine.

30

As shown in Figure 10, system board 170 is connected directly to each transceiver in the network node through conductors in a cable 175, one cable provided for each transceiver. Each cable 175 consists of one pair of conductors for carrying data signals from switching engine 182 to a transceiver, one pair of

conductors for carrying data signals from the transceiver to the switching engine, one conductor for carrying serial data from CPU 173 to a micro controller on the transceiver board (described below) and one conductor for carrying serial data from the transceiver micro controller back to CPU 173. Cable 175 includes
5 additional cables for providing power and ground to the transceiver. As described here, CPU 173 communicates with each transceiver micro controller by means of serial data.

Processor 173 is also connected to the access device 172 by means of a cable
10 through Fast Ethernet (100-Base TX) physical layer interface (PHY) device 178 and Ethernet connector 171. The cable, commonly known in the industry as Category -5 cable, consists of 8 individual conductors and is widely used to carry Fast Ethernet (100-Base TX) data. Following the industry standard, four of the conductors are used to carry 100-Base-T data signals between access device
15 172 and CPU 173. The remainder of the 4 conductors is used to provide unregulated 48Volts DC power and ground from the access device to the switching power supply 188. Both ends of the cable are terminated using the industry standard connectors commonly referred to as RJ-45. The unregulated 48 Volts DC power is stepped down by the switching power supply 188 to provide
20 regulated power supplies to the system board as required by various electronics components in the system. Power supply 188 also provides powers to the transceivers boards 220A-D through cable 175.

In one exemplary embodiment, processor 173 comprises an AMD brand
25 Au1000N processor and switching engine 182 comprises a Galileo brand GT-48511A Fast Ethernet/IP switching engine. It will be appreciated that other processors may be substituted for CPU 173 and other packet switching devices for switching engine 182.

30 One exemplary user computer 190 is shown, for example comprising a typical personal computer, connected to access device 172 through a high-speed digital connection, for example an industry standard Ethernet Cat 5 cable connection. User computer 190 is, for example, contained in one of the neighborhood homes described above. A conventional modem 194 is further provided on system

board 170 and connected to the Internet 103. Modem 194 is used in the manner described below to access a network system server 202 for initial setup and various other local processes, the network system server used to control the mesh IP network in the manner described below. Modem 194 can be used, for
5 example, to dial up system server 202 directly or through an intermediate ISP. Alternatively, a short-range wireless radio frequency network communications card (not shown) can be built into each network node, the modem and/or wireless radio frequency network communication card enabling adjacent nodes to communicate during, for example, installation and alignment procedures of
10 the type described below.

Switching engine 182 performs the primary function of switching and routing high-speed IP data packets. In the embodiment described herein, switching engine 182 receives IP data packets encapsulated as Fast Ethernet data packets
15 from the six interfaces: the four 100-Base FX (Fast Ethernet) PHYs 186A-D connected to the transceivers, the one 100-Base TX (Fast Ethernet) interface PHYs 180, 181 connected to the CPU 173 and the one bridge interface 185 connected to the CPU.

20 The switching engine 182 makes the switching decision on each packet based on information contained within its data packet header, such as the IP header. As is known in the art, the IP header consists of the IP addresses of the sender and the destination of the packet. For added flexibility in making a routing decision, the switching engine can also use information contained within a Fast Ethernet
25 packet header. As is also known in the art, a Fast Ethernet packet header consists of, among other things, source address, destination address, VLAN ID etc. In order to provide more sophisticated network services to each packet, the switching engine may also examine other parameters contained within each IP packet, such as the TCP port number. Once the data packet switching engine 182
30 has processed the incoming information and made a data routing decision, the data packet can be transmitted out on any one of its six interfaces.

In the present embodiment of the invention, processor 173 serves three primary functions. It serves, through the switching engine 182, as an interface between

the mesh network of infrared nodes and the user accessing the network through access device 172. This function allows for network operators to implement any kind of processing of packets received from each user before such packets enter the network through switching engine 182. Such processing includes but is not
5 limited to authentication, encryption, data rate limitation, etc.

The processor further serves as an agent of network system server 202 for configuration and management of switching engine 182. Such configuration and management may include but are not limited to managing the switching engine's
10 routing table update and network failure recovery.

Processor 173 further performs tracking functions between infrared network nodes, descriptions of which are detailed below. In the described embodiment, CPU 173 operates using the Linux operating system supporting sub-processes
15 needed to implement functions described above. Alternate operating systems, for example Windows CE or others, would likewise suffice.

With reference now to Figure 11, one exemplary embodiment of access device 172 is shown including a power supply 172A and a signal conductor 172B.
20 Power supply 172A includes a conventional 60Hz transformer 210 connected serially to a conventional rectifier and capacitor circuit 212 and a 1 Amp limiter 214 for converting 120Volt AC to unregulated 48Volt DC. Signal conductor 172B includes an appropriately connected, grounded Ethernet connector 216 for connecting to user computer 190. The 48Volt output of power supply 172A is
25 used to provide power to Ethernet connector 218 and hence to system board 170.

With reference now to Figures 12 and 13, there is shown a single transceiver electronics board 220A, mounted on transceiver mount 115A (see Figs. 2, 3) connected to the system board 170 (Figure 10) via power conductor 189 from
30 power supply 188 (Figure 10) and cable 175. For each transceiver board such as 220A, cable 175 includes two lines of control data from processor 173 for controlling the elevation and azimuth stepper motors 111A, 113A, respectively, and four lines of Ethernet data for transmitting and receiving data.

- Transceiver electronics board 220A supports power and data conductors, indicated at power/data connector 222, the power being connected to all the transceiver sub-systems, the control data signals being connected to a microprocessor 224, and the Ethernet data signals being connected to a power amplifier 226 and a transimpedance amplifier 228. An elevation motor connector 230 is connected between microprocessor 224 and elevation motor 111A. An azimuth motor connector 232 is likewise connected between microprocessor 224 and azimuth motor 113A.
- 5
- 10 Elevation and azimuth motors 111A and 113A, respectively, are shown connected to transceiver base 117A, the base and mount 115A connected by two separate 14:1 gear assemblies for dividing down the arcuate motion of the motors as described herein above.
- 15 LED 122A, diffuser sheet 120A, dome lens 116A and APD 118A, each mounted separately on transceiver mount 115A in the optics assembly described herein above, are for explanatory purposes shown in Figure 12 in dotted line. LED 122A and APD 118A are additionally shown schematically in the schematic diagram of Figure 13.
- 20 A signal strength indicator 227 is connected to microprocessor 224 for sensing the strength of an incoming signal detected by APD 118A through transimpedance amplifier 228.
- 25 The control signal data from the conductors in cable 175 is routed through power/data connector 222 to microprocessor 224, the microprocessor in turn providing control signals to the elevation and azimuth stepper motors through the connectors 230, 232.
- 30 The Ethernet data from the appropriate conductors in cable 175 includes 2 transmit data conductors connected to power amplifier 226 for driving LED 122A to transmit Ethernet data and 2 receive data connectors connected to transimpedance amplifier 128 for receiving Ethernet data detected by APD 118A.

With reference now to Figures 14A and B, transimpedance amplifier 228 is seen, in accordance with one embodiment of the present invention, to include three series-connected differential pair amplifiers (DPAs) 250, 252, 254, each
5 consisting of two NPN RF transistors. APD 118A is connected to the inputs of DPA 250 across a resistor 256 through a pair of capacitors 258, 260, capacitor 258 connected to the positive input of DPA 250 while capacitor 260 is connected to the negative input of the DPA. A resistor 262 is connected between the positive output of DPA 254 and the negative input of 250. A resistor 264 is
10 connected between the negative output of DPA 252 and the positive input of DPA 250.

In operation first with respect to transimpedance amplifier 228, the amplifier performs the well-known function of converting an input current to a limited or
15 clipped output voltage, whereby to amplify the relatively low photo-detective current generated by APD 118 upon the receipt of IR signals. The design using feedback resistors 262, 264 provides outstanding gain and sensitivity without the need for expensive matched-pair differential amplifiers typical of the prior art. The frequency response of the amplifier can also be well controlled, eliminating
20 the need for additional filtering.

In operation with respect to the transceiver electronics board 220A and its interaction with system board 170 (Figure 10), the boards provide a plurality of functions.

25 More specifically, with respect to the elevation and azimuth position of the transceiver optics, transceiver electronics board 220A functions to receive motor control signals from CPU 173 (Figure 10) and provide them via microprocessor 224 and elevation and azimuth motor connectors 230, 232 to the respective
30 elevation and azimuth stepper motors 111A, 113A.

With respect to data transmission, transceiver electronics board 220A functions to receive Ethernet data from PHYs 186A, B, C, D (see Figure 10) through connector 222 to drive LED 122A through power amplifier 226 for transmitting

IR data to an adjoining node (see Figure 1). Data from the user computer is conveyed through various intervening connectors, conductors and converters through CPU 173 to switching engine 182 and data from other transceivers are relayed through switching engine 182. As described elsewhere herein, in the present embodiment the network comprises an IP network, managing packet-switched data in accordance with Internet protocol standards.

With respect to data receipt, transimpedance amplifier 228 converts current generated through the receipt of IR data by APD 118A into electronic signals, which are transmitted through connector 222 to PHYs 186A, B, C, D. Received data which is destined for the user computer is passed from switching engine 182 through CPU 173 and the various intervening connectors, conductors and converters, while data from other transceivers are relayed through switching engine 182 for transmission by another transceiver to another node.

With respect to the installation, alignment and tracking of the transceivers in the system node, signal strength indicator 227 detects the relative power of the incoming signal received by APD 118A into transimpedance amplifier 228 and transmits the same to CPU 173. This signal strength is used in the manners described below to align newly installed network node transceivers and to realign existing transceivers already on the network.

It will thus be seen from a consideration of Figures 10, 11 and 12 that system board 170, of which there are one per network node 106, contains electronics pertinent to each of the transceivers within the node. The transceiver electronics boards, of which there is one per transceiver 108A-D, contain electronics pertinent to the operation of each individual transceiver.

From a consideration of the above, it will be understood that high-level IP network management functions are determined by network system server 202 and communicated to the various system nodes for storage in the system and transceiver boards. Local control of the nodes and transceivers is performed through the operation of the system and transceiver boards in each node. It will be further understood that certain operating programs and control information

may be loaded onto the system and transceiver boards upon assembly so that it is available before the installation of the node into the network, for example to facilitate the initial tracking process described below. It will be apparent that many different strategies for loading and updating data and software within the nodes may be implemented in accordance with the present invention.

While the present embodiment of the invention has been shown and described with respect to one neighborhood network of nodes, it will be understood that multiple such networks can exist over greatly diverse geographical areas. Such networks can communicate and exchange data with each other through intermediary networks. In one embodiment, for example, multiple neighborhood networks may connect through a common system node(s) 202, which provides common management support to each neighborhood network. In another embodiment, various neighborhood networks may be separately managed through separate network servers and connect through one or more intermediary servers and or intermediary networks of differing type. It will now be apparent that numerous configurations of networks can exist in accordance with the present invention.

20 Network Operation – New Node Installation

With reference now to Figures 15A & B there is shown a process 300 for installing a new node 106 into an established mesh network of nodes 100 (Figure 1). As described, a tracking or synchronization process is initiated whereby selected transceivers within the new and existing nodes are relatively positioned such that those selected transceivers can exchange data. The new node is then acquired into the existing network, becoming an active node within the network.

Initially, new node 106 is installed on an upper, outer surface of a structure convenient to a house 102 or other structure containing a computer to be connected to the network (step 302). The node may be installed on the wall, roof or chimney area of a house or on an adjoining structure such as a lamp or utility post. The exact position of the node is selected to provide a line-of-sight to at least one existing node located within the network and within communicating distance. The new node is preferably connected to a user

computer within the house through the node's Ethernet connection. This enables the node to communicate with the user both during the installation procedure and subsequently during data transmission and receipt. Every new node is installed level, that is with a zero degree angle of elevation.

5

A homeowner or a professional installer may install the new node. It will be appreciated that, in accordance with the invention, the installation of the new node is typically sufficiently simple to enable non-professional installation.

- 10 Upon securing new node 106 to the supporting structure, the installer orients the node by twisting the entire node assembly so that a particular mark on the outer surface of the weatherproof housing 112 is pointed to compass north (step 304). With the transceivers pre-positioned at the factory in the manner described above, the relative orientation of each transceiver within the node is known,
- 15 within an error margin, to network server 202. The newly installed node's location is communicated to network server 202 (Figure 10) (step 306), for example through a dial-up modem connection on the user computer, or by a telephone call to the server placed by the installer, or using modem 194 on system board 170 (Figure 10). Alternatively, this communication may be made
- 20 through the wireless RF network described above.

Network server 202 examines a stored database of nodes within the local network (step 308) to determine adjoining nodes likely to have line-of-sight with the newly installed node (step 310).

25

- If there is no line-of-sight network node contained in the database, the installer, typically a professional installer in this circumstance, will provide the network server 202 with a specific position of an existing network node with which the new node can be synchronized (step 314). This can be done using one of the
- 30 many communication options described herein above. If one or more line-of-sight nodes exist in the network, the network server will select one existing node for synchronization with the new node (step 316). If more than one line-of-sight nodes exist in the network, a single node is selected using criteria, for example, based on expected load and usage of the various nodes.

To facilitate the actual acquisition of the new node into the network, network server 202 determines necessary synchronization data for both nodes, including but not limited to: the likely best transceivers to use (it will be recalled from a
5 consideration of the above that each node contains four transceivers), a first best guess of the initial directional orientation of each transceiver likely to result in line-of-sight communication following the process described below, the start time to begin the process of synchronization for acquiring the new transceiver into the network, the current time and the angular sweep range of the new and
10 existing transceivers (step 318).

It will be appreciated that the existing transceiver, being integrated into the network, has an accurate determination of a given reference orientation. In contrast, the new transceiver is only approximately oriented to a reference
15 direction by the installer and may have an error in orientation. As will be described in further detail below, during the initial synchronization both the new and existing nodes are swept through the initially determined sweep range likely to insure line-of-sight communication, with the new node additionally sweeping through the angle of uncertainty whereby to account for that uncertainty.
20

Upon completing the calculation and determination of the relevant synchronization data, server 202 transmits the necessary data to the existing node (step 320) and to the new node (step 324), where the data is stored for the subsequent synchronization and acquisition. Data is transmitted to the existing
25 node through the network. As described above, even though the new node is not active in the network and able to receive network data, synchronization data can be provided by one of the many available options described above, i.e. through the built-in modem, through an upload from the user computer, or through a wireless RF connection.
30

Network Operation – New Node Synchronization and Acquisition

With reference to Figures 16A&B, there is shown a process 330 for synchronizing a new transceiver in a new node with an existing transceiver in an existing node to acquire the new node into the network.

Initially, a clock calibration occurs between the new and existing nodes, using the described alternate means of communication, prior to initiating the described synchronization process. The nodes may further agree to the pre-established or a
5 different start time.

To begin the synchronization, the new and existing nodes each move the selected active transceiver to its respective first best guess of orientation (step 332), each transceiver being maintained level. This is accomplished using the above-
10 described azimuth motors to rotate the transceivers to the starting positions. At the synchronous start time, the azimuth motors are operated within each of the new and existing transceivers whereby to actuate the sweeps within the initially determined angular sweep ranges, both the existing and new transceivers sweeping through the sweep range at approximately the same average speed,
15 with the new node sweeping in an additional back and forth pattern having an angular width of the allowed reference orientation uncertainty (step 334). During the sweeps, relative signal strengths of detected infrared light for each position in the sweep pattern are stored within each of the new and existing nodes.

20 If a signal is not detected by either transceiver (step 336), then the elevation angle of each transceiver is adjusted in a complementary manner in selected beam-width increments and the sweeps are repeated (step 331). For example, when one-half beam-width increments are selected, the new node will adjust
25 down in elevation approximately one-half the estimated beam width while the existing node will adjust up in elevation approximately one-half the estimated beam width, and the above-described angular sweep repeats. This process of elevation adjustment and sweep is repeated as long as no signal is detected (step 333) by either transceiver and the final limits of both devices are not reached
30 (step 335). Each time a sweep fails to result in a transceiver detection for a particular elevation (step 333) and the device limits are not reached (step 335), that same sweep is repeated in the opposite azimuth direction (step 337) before the elevation angle is again adjusted. It will be understood that various beam-width increments can be used to accomplish the alignment results.

If, after completing sweeps in both directions at the final elevation limits of both transceivers (step 335) infrared light signals are not detected by either the new or existing transceiver, a failure signal is generated by the existing transceiver (step 5 338) and transmitted to the system board CPU and/or the network server 202 (step 340). The CPU and/or the network server then selects a new existing transceiver (step 342) and provides the tracking data to this existing node, as described above, to restart the synchronization process (step 344) as described above (steps 332-337).

10

The next existing node may be selected in real-time by the network server or may have been earlier identified by the network server and stored in the system board CPU for use in the event that communications couldn't be established with the first-choice existing node. If, upon repeating the above-described sweep 15 process neither transceiver can detect an infrared signal from each other, a failure indicator is again transmitted to the network server which determines that the new node is faulty and must be reinstalled or replaced (step 346).

If either the new or existing transceiver receives signals during any one sweep 20 (steps 333 or 336) then, upon the completion of that sweep pattern, both transceivers return to the position of maximum signal strength (step 348). It will be understood that, due to the symmetry of the co-aligned transceiver optical systems described above, if either transceiver detects a signal, the assumption is made that the other transceiver has also detected a signal and that both 25 transceivers are returning to the position of maximum infrared light signal strength.

Following the completion of the sweep pattern, the new and existing transceivers execute asynchronous tracking events (described in detail below) (steps 350, 30 352) one or more times (step 354), followed by an attempt to communicate actual signal data (step 356).

If signal data communication cannot be established within a predetermined number of asynchronous tracking events (step 358) as determined by the system

operator, then the sweep pattern is reinitiated (step 334) to re-determine the position of maximum signal (step 348) and the asynchronous tracking events (steps 350, 352) begin again. When signal data communication is established (step 358), then the synchronization of the new and existing transceivers is

5 complete with the new node acquired into the existing network (step 360). The new node is now installed into the network and can function as a normal network node including communicating network data and cooperating to install new nodes into the network. It will be understood that, once installed into the network, the actual orientation of the newly installed node is known and stored

10 for use in subsequent tracking activities. The user of the newly installed node is thus connected to the Internet (or other network.)

Network Operation – Asynchronous Tracking Event

Immediately following the install process described above and periodically

15 thereafter it is necessary to cause a communicating transceiver pair within a network to change its physical position so as to establish and maintain good optical alignment and signal communications between the adjoining transceivers. This process, initiated by network server 202 or by an individual transceiver, is termed a 'tracking event' and is described with respect to the Figure 17 process

20 illustrating an asynchronous tracking event 370.

The transceivers are tracked asynchronously, each transceiver beginning its tracking event upon the initiation of the network server or on the occurrence of a predetermined time (step 372). Upon initial installation, as described above, an

25 asynchronous tracking event is initiated substantially immediately upon both transceivers returning to their respective positions of maximum signal strength. Upon lost or diminished signal, as described below, an asynchronous tracking event occurs within a certain time slot for each transceiver.

30 Initially, the relative distance between the two nodes is determined by the node location data contained in network server 202 (step 374). If the distance is less than a predetermined distance, typically about 10 meters, then processor 173 on system board 170 (Figure 10) of the node containing the transceiver to be repositioned initiates a transceiver centroid-tracking path (step 382). That is, the

respective elevation and azimuth motors are operated to move the transceiver in a centroid path, simultaneously collecting and storing signal strength to determine the position yielding maximum signal strength (step 384). At the completion of the centroid tracking path, the transceiver is controlled to return to
5 the position of maximum received signal strength (step 386).

If the actual distance is greater than the predetermined distance (step 374), the transceiver is similarly controlled as above but in a dither tracking path (step 376). Again, receive signal strength is recorded to determine a maximum (step
10 378) to which the transceiver is returned at the end of the dither path (step 380).

As will be understood from the description of the new node installation process 330 described above, in the asynchronous tracking process, each of the two transceivers involved in the process alternate motion in the tracking event. That
15 is, the first transceiver will complete a tracking path and return to the position of maximum signal strength. Next, the second transceiver will complete a tracking path and likewise return to the position of maximum signal strength. The process of alternating transceivers for tracking events continues for a predetermined number of times. This predetermined number of times may
20 depend on the occurrence of an event, such as the ability to communicate data, or on the completion of a predetermined, stored number of tracking events.

It will be seen from a consideration of the acquisition process described above in combination with the asynchronous tracking event process, that upon the
25 installation of a new node a simple, simultaneous or synchronous sweep in a circular, horizontal path is used to determine initial transceiver positioning, while the appropriate centroid or dither tracking path is used in an asynchronous process to fine-tune transceiver position for maximum signal strength.

30 It will be appreciated that numerous other tracking paths may be selected to determine optimum transceiver positioning.

Network Operation – Unscheduled Tracking Event

Periodically, due to a variety of conditions including but not limited to changes in tolerance of mechanical components, changes in supporting structures and underlying geography and changes in weather, a transceiver within a network
5 node will lose good signal communications with its 'mate' in the adjoining node. When this happens an unscheduled tracking event process 400 initiates as is shown in Figure 18.

Upon assembly and initial programming, each transceiver in each node is given a
10 pre-assigned default time slot within which to initiate a tracking event upon the occurrence of an unscheduled signal loss. This default time slot may, for example, be stored on system board 170 in memory 176. No two transceivers within the same neighborhood are provided with the same pre-assigned default time slot. It will be understood that, in this manner, the asynchronous tracking
15 process 370 (Figure 17) may be performed automatically and without need for the network server to communicate with each transceiver.

Upon the detection of minimally acceptable receive signal strength (step 402), the transceivers to be tracked each begin an asynchronous tracking event within
20 their respective pre-assigned time range (step 404). The asynchronous tracking is performed in accordance with Figure 17 described above, and if signals are detected and communication is established (step 406) the process terminates (step 408). If no signal is detected by the transceivers and no communication is established, then a service event is initiated (step 410). A service event may, for
25 example, include the transmission of a service notice to a user or a modem communication to the network server. A service event may, for example, result in human intervention to replace a defective network node.

Network Operation – Digital Signal Quality-Based Acquisition, Alignment and 30 Tracking

It will be understood that, in the above-described embodiment, acquisition, alignment and tracking of the transceivers are achieved using the detected signal strength of the IR light beam. In this alternate embodiment of the invention, acquisition, alignment and tracking are achieved by testing the quality of digital

signal packets transmitted between nodes. In this embodiment, error tracking is performed based on packets received from the far side transceiver, i.e. the transceiver being pointed at. The quality of the digital signal test packets sent by the far-side transceiver is used as an indicator of signal quality and tracking is performed based on that metric. By receiving only packets with the matching source and destination MAC addresses, even retro-reflection will not diminish the quality metric.

One problem associated with measuring link quality is the time required to get a meaningful error rate from a small amount of traffic. For example, when sending twenty 1440-byte packets each second, the following parameters, including bit error rate (BER), show the approximate amount of time required to measure a particular error rate to within a 75% confidence interval.

15	Percent	BER	TIME
	50%	250×10^{-7}	0.8s
	90%	50×10^{-7}	4s
	99%	10×10^{-7}	20s
	99.5%	5×10^{-7}	40s
20	99.9%	1×10^{-7}	200s

Generation of twice as much traffic cuts the confidence interval in half. However, because generation of test traffic is CPU intensive, there is a limit to the amount of traffic that a node can generate. However, in the context of tracking, speed of measurement is most important on poor signals and least important on good signals. That is, it is faster to determine an acceptable error rate within a certain level of confidence than it is to determine an exact error rate. To take advantage of this, the error-rate determination algorithm described herein counts only enough errors, for example, 4 errors, to make the error-rate determination within a 75% confidence interval. If no packets are received, the algorithm assumes after one second that four packets were lost and the error rate can be determined.

In this embodiment of the invention, the test packets sent from the transmitting transceivers contain the destination and source MAC address of the node processor so that the packets can be received and processed. Contained within each packet is a sequential counter which is used for counting lost packets and a tracking event number which is used to prevent tracking conflicts. The payload and length of the packets are alternated every other packet between two packet types: a Long Packet and a Short Packet.

A Long Packet comprises a 1440-byte packet. The payload can be set to a challenging frequency distribution to bring out errors earlier than a random frequency distribution packet would.

A Short Packet comprises a 64-byte packet consisting primarily of random data (the counters). This packet is easy to transmit and makes possible measurement of marginal links.

The combination of these packets allows measurement of the error rate over a large range of qualities without the requirement of sending as much data as would be required with large random packets. This allows greater control when tracking, speeds the measurement of error rate, and reduces processor load.

The sequential counter is 16-bits unsigned, and is encoded as a hexadecimal text value. If a node is in the process of tracking, a random 16-bit hexadecimal value will be added to the test packet on the tracking transceiver. In the event that two transceivers begin to simultaneously execute tracking processes, this hexadecimal value is used to prevent tracking conflicts by forcing the transceiver with the lower random value to cease tracking. The transceiver with the higher random value is permitted to complete the tracking process.

In accordance with this embodiment of the invention, one method of tracking comprises the process of taking a predetermined number of steps in one direction and then the same number of predetermined steps in the opposite direction. As described above, each 'step' comprises a step of the controlling transceiver stepper motor, which in turn corresponds to an incrementally changed position

of the controlled transceiver. At each stepped position, as well as at center, the error rate is measured. If one position is better than then other two, the process optimizes transceiver positioning by moving the transceiver to the position of the best detected digital signal quality, i.e. the position of least error rate. If the two
5 best positions are equally good, the process splits the difference in position between the two. Three examples are shown below.

Example 1:

99.5% at 4 steps left
10 98.0% at center
94.3% at 4 steps right
MOVE 4 STEPS LEFT

Example 2:

15 99% at 4 steps left
100% at center
100% at 4 steps right
MOVE 2 STEPS RIGHT

20 Example 3

99% at 4 steps left
99.5% at center
88% at 4 steps right
DO NOTHING

25

The examples above all use moves of four steps left or right (i.e. azimuth), but this stepping, or vectoring, applies equally to vertical (i.e. elevation) and diagonal tracking movements. The number of steps attempted depends on the size of the IR spot as a function of range and whether the tracking process is in a
30 phase of rough centration, fine tuning, or searching for a signal.

With respect to the alignment process described here, references to the IR spot size comprise the number of steps over which a good quality digital signal can be obtained by the receiving transceiver under alignment. Due to the shape of

the beam, at shorter ranges, that is distance between transceivers, a good quality signal can be received over a wider angle or larger number of steps. As range increases, only the central, highest intensity, portion of the beam yields a good digital signal. Thus, the size of the beam determines the angle to be swept in
5 determining the edges of good link quality.

One full tracking sequence is described below based on the alignment process described above. The number of steps N is based on typical spot size at the particular range.

10

With reference now to Figure 19, a tracking process based on digital signal quality is shown wherein the best digital signal quality is measured and stored at the current starting position (step 502), that is, without moving the transceiver to be aligned. As noted above, it will be understood that the best signal quality
15 corresponds to the lowest error rate. The transceiver being positioned is then stepped in elevation so as to measure the quality N steps up and N steps down from the center starting position, and moved based on the vectoring process described above (step 504), i.e. to the strongest signal position or, if two signals are of equal strength, to a position half-way between them. The position
20 yielding the best quality signal as determined from the prior up/down measuring and repositioning steps is used as the new starting position, with the best quality signal used as the new best quality signal (step 506).

The transceiver being positioned is then stepped in azimuth so as to measure the
25 quality N steps left and N steps right from the new starting position, and again moved based on the vectoring process described above (step 508), i.e. to the strongest signal position or, if two signals are of equal strength, to a position half-way between them. The position yielding the best quality signal as determined from these left/right measuring and repositioning steps is used as the
30 new starting position, with the best quality signal used as the new best quality signal (step 510).

The transceiver being positioned is then stepped in elevation so as to measure the quality 1 step up and 1 step down from the new starting position, and again

5 moved based on the vectoring process described above (step 512), i.e. to the strongest signal position or, if two signals are of equal strength, to a position half-way between them. If a position adjustment is made based on signal quality, the resulting best signal quality and position become the new best signal quality and position.

10 The transceiver is then stepped in azimuth so as to measure the quality 1 step left and 1 step right from the new starting position, and again moved based on the vectoring process described above (step 514), i.e. to the strongest signal position or, if two signals are of equal strength, to a position half-way between them.

15 Optionally, the transceiver can be similarly stepped and adjusted along diagonals through repetition of the stepping/measuring/new starting position, steps described above, the only change being that the transceiver is stepped along diagonal lines relative to its starting position. It will also be apparent that other tracks alternative to up, down, left, right and diagonals can be selected for which to repeat the step-measure-set new starting point, process described above

20 This full tracking sequence as described with respect to Figure 19 achieves gross pointing and edge avoidance as well as fine tuning.

25 Between tracking events, the receive quality is constantly measured by counting the test packets received. To avoid spurious moves based on temporary drops in quality, at the initiation of a tracking process, in lieu of measuring a new center quality value, the highest of the previous three quality values are taken as the center quality measure. If a move is performed in a subsequent stage of tracking, the quality at the new point is taken as the new quality at center.

30 When a link is interrupted by an environmental event like fog or a blockage, it's possible for transceivers to become significantly misaligned during the blocked period. In order to track back onto a spot after the blockage is removed, a transceiver may have to move more steps than would be required for normal tracking. In this recovery mode, the transceiver moves left, right, up, and down as well as all diagonals. The number of steps moved from center in each

direction is 0.5S, 1.0S, and 1.5S where S is the number of steps that the spot is wide.

5 This recovery algorithm is enough to bring a transceiver back to alignment as long as one side of a link is pointed such that its transmit beam is nearly on target. In an alternate embodiment of this recovery process, the transceiver being recovered is returned to the last known good position and alignment as described above is performed about that point.

10 In order to provide some flexibility in the behavior of tracking, in one embodiment of the invention a transceiver is modeled as being in one of three states: good quality, poor quality, and no signal

15 The tracking frequency can be set differently for each of these states. In one example, the "no signal" state is defined as having no test packets for 15 minutes. The dividing line between "good quality" and "poor quality" is configurable but should typically be between 80% and 99% of test packets arriving without error. Four parameters which affect tracking and which are controllable for each transceiver are listed below.

20

Nominal Quality Threshold: The percentage of good test packets received which are determinative between "good quality" and "poor quality" states.

Good Tracking Frequency: How frequently in time a transceiver should track when quality is good.

25 Poor Tracking Frequency: How frequently in time a transceiver should track when quality is poor.

Recovery Tracking Frequency: How frequently in time a transceiver should track in recovery mode, when there is no signal.

30 The nominal quality threshold is an integer representing percentage and the various tracking frequencies are integers in minutes. A tracking frequency of less than 15 minutes may be treated as a 15 minute tracking frequency. A tracking frequency of 0 indicates that tracking should never be done in that state.

While the above-described tracking processes have been illustrated with respect to infrared transceivers, it will be understood that they similarly apply to other point-to-point communications systems, for example including narrow-band radio frequency and microwave transmitters/receivers, all included herein as point-to-point communications systems or transceivers.

There is thus provided new and improved infrared transceivers, infrared network nodes comprising multiple transceivers, an infrared mesh network, various circuits and various processes for installing and synchronizing new nodes whereby to extend existing high-speed digital network capacity into areas inaccessible by conventional networks for the reasons described above.

The transceiver transmitter optics are designed using infrared light with a relatively wide beam divergence in the range of about 0.5-1.0 degrees. In combination with sensitive receive optics having a slightly smaller field of view, much flexibility is provided in the placement and alignment of the transceivers. This makes installation and expansion of the network cost-effective and straightforward.

Automated install and synchronization processes operate to keep the network functional and to simplify new installations and changes in network node location.

The system has commercial application in the field of high-speed digital data networks including Internet networks such as those used by Internet Service Providers.